
Bi-level Score Matching for Learning Energy-based Latent Variable Models

Fan Bao*, Chongxuan Li*, Kun Xu, Hang Su[†], Jun Zhu[†], Bo Zhang
Dept. of Comp. Sci. & Tech., Institute for AI, THBI Lab, BNRist Center,
State Key Lab for Intell. Tech. & Sys., Tsinghua University, Beijing, China
bf19@mails.tsinghua.edu.cn, {chongxuanli1991, kunxu.thu}@gmail.com,
{suhangss, dcszj, dcszb}@tsinghua.edu.cn

Abstract

Score matching (SM) [24] provides a compelling approach to learn energy-based models (EBMs) by avoiding the calculation of partition function. However, it remains largely open to learn energy-based latent variable models (EBLVMs), except some special cases. This paper presents a bi-level score matching (BiSM) method to learn EBLVMs with general structures by reformulating SM as a bi-level optimization problem. The higher level introduces a variational posterior of the latent variables and optimizes a modified SM objective, and the lower level optimizes the variational posterior to fit the true posterior. To solve BiSM efficiently, we develop a stochastic optimization algorithm with gradient unrolling. Theoretically, we analyze the consistency of BiSM and the convergence of the stochastic algorithm. Empirically, we show the promise of BiSM in Gaussian restricted Boltzmann machines and highly nonstructural EBLVMs parameterized by deep convolutional neural networks. BiSM is comparable to the widely adopted contrastive divergence and SM methods when they are applicable; and can learn complex EBLVMs with intractable posteriors to generate natural images.

1 Introduction

An energy-based model (EBM) [35] employs an energy function mapping a configuration of variables to a scalar to define a Gibbs distribution, whose density is proportional to the exponential negative energy. Being flexible, EBMs can naturally incorporate latent variables to fit complex data and extract features. Among them, representative models including restricted Boltzmann machines (RBMs) [21], deep belief networks (DBNs) [23] and deep Boltzmann machines (DBMs) [48] have been widely adopted [60, 55]. However, it is challenging to learn EBMs because of the presence of the partition function, which is an integral over all possible configurations, especially when latent variables present.

The most widely used training approach is the maximum likelihood estimate (MLE), or equivalently minimizing the KL divergence. Such methods often adopt Markov chain Monte Carlo (MCMC) [42] or variational inference (VI) [29] to estimate the partition function (or its gradient with respect to the model parameters). Contrastive divergence (CD) [21] and its variants [57, 63, 43, 44, 9, 10, 16] are proven effective in models with fully visible variables or tractable posteriors of latent variables (e.g., RBMs). The recent work [63] present an unbiased version of contrastive divergence. Besides, Ingraham and Marks [26] perform approximate Bayesian inference over EBMs. In deep models such as DBNs and DBMs, previous work [23, 36, 48] often adopts layer-wise training strategy. Recently, several methods [34, 37] attempt to learn general energy-based latent variable models (EBLVMs) in a black-box manner by VI. In these methods, the problem of inferring the latent variables is addressed

*Equal contribution. [†] Corresponding author.

by advances in amortized inference [31] but the variational bounds for the partition function are either of high-bias [37] or high-variance [34] on high-dimensional data.

Score matching (SM) [24] provides a promising alternative approach to learning EBMs. Compared with MLE, SM does not need to access the partition function because of its foundation on Fisher divergence minimization [28], while involves a second order derivative. Many attempts [32, 59, 51, 54, 38, 45] try to estimate the second order derivative efficiently and recent work [38] can scale up to natural images. However, it is much more challenging to incorporate latent variables in SM than in MLE because of its specific form. As far as we know, extensions of SM for EBLVMs [56, 58] make strong structural assumptions that the posterior is tractable [56] or in the exponential family [58].

Considering the complementary advantages between MLE and SM, a natural question arises:

Can we infer the latent variables in nonstructural EBLVMs using amortized inference as in MLE and at the same time learn such models without explicitly estimating the partition function as in SM?

In this paper, we present bi-level score matching (BiSM), which is generally applicable to existing SM objectives [59, 54, 38] to learn EBLVMs with a minimal model assumption. The key to our approach is to reformulate a given SM objective as a bi-level optimization problem. The higher level problem modifies the original SM objective by approximating the marginal model distribution with the ratio of the model distribution over a variational posterior. The lower level problem optimizes certain divergence between the variational posterior and the true one. By reformulating the divergence used in the lower level problem, BiSM only needs to access the model energy and the variational posterior in its calculation. Further, under the nonparametric assumption [15], BiSM is equivalent to the original SM objective (see Theorem 1). To solve BiSM efficiently, we propose a practical algorithm using alternative stochastic gradient descent with gradient unrolling [41] and formally characterize the gradient bias and the convergence of the algorithm (see Theorem 2 and Corollary 3).

We evaluate BiSM on two EBLVMs. The first model is the well-known Gaussian restricted Boltzmann machine (GRBM) [60, 22]. We compare BiSM with the corresponding SM methods [56, 54] and the contrastive divergence (CD) [21, 57]. On a toy 2-D dataset and the Frey face dataset, BiSM achieves comparable performance to these strong baselines. The second model is a highly nonstructural EBLVM parameterized by deep convolutional neural networks to fit natural images. The CD and SM-based baselines are not applicable because of the intractable posteriors while BiSM can perform inference and learning successfully. We show the promise of BiSM by testing the sample quality and the inference accuracy on the MNIST and CIFAR10 datasets. To the best of our knowledge, previous state-of-the-art EBLVMs [23, 36, 48] cannot generate natural images in such a purely unsupervised learning setting.

2 Preliminaries

In this section, we present preliminaries about the Fisher divergence [28] and existing score matching methods in energy-based models (EBMs). Formally, an EBM defines a distribution: $p(\mathbf{w}; \boldsymbol{\theta}) = \tilde{p}(\mathbf{w}; \boldsymbol{\theta}) / \mathcal{Z}(\boldsymbol{\theta}) = e^{-\mathcal{E}(\mathbf{w}; \boldsymbol{\theta})} / \mathcal{Z}(\boldsymbol{\theta})$, where $\mathcal{E}(\mathbf{w}; \boldsymbol{\theta})$ is the associated energy function parameterized by learnable parameters $\boldsymbol{\theta}$, $\tilde{p}(\mathbf{w}; \boldsymbol{\theta})$ is the unnormalized density, and $\mathcal{Z}(\boldsymbol{\theta}) = \int e^{-\mathcal{E}(\mathbf{w}; \boldsymbol{\theta})} d\mathbf{w}$ is the partition function. Here, we assume that the variable \mathbf{w} is fully visible and continuous.

Fisher divergence The Fisher divergence [28] between the empirical data distribution $p_D(\mathbf{w})$ and the model distribution $p(\mathbf{w}; \boldsymbol{\theta})$ is defined as:

$$\mathcal{D}_F(p_D(\mathbf{w}) || p(\mathbf{w}; \boldsymbol{\theta})) \triangleq \frac{1}{2} \mathbb{E}_{p_D(\mathbf{w})} [||\nabla_{\mathbf{w}} \log p(\mathbf{w}; \boldsymbol{\theta}) - \nabla_{\mathbf{w}} \log p_D(\mathbf{w})||_2^2], \quad (1)$$

where $\nabla_{\mathbf{w}} \log p(\mathbf{w}; \boldsymbol{\theta})$ and $\nabla_{\mathbf{w}} \log p_D(\mathbf{w})$ are the model score function and data score function [24], respectively. The model score function does not depend on the value of $\mathcal{Z}(\boldsymbol{\theta})$. Indeed, we have:

$$\nabla_{\mathbf{w}} \log p(\mathbf{w}; \boldsymbol{\theta}) = \nabla_{\mathbf{w}} \log \tilde{p}(\mathbf{w}; \boldsymbol{\theta}) - \nabla_{\mathbf{w}} \log \mathcal{Z}(\boldsymbol{\theta}) = \nabla_{\mathbf{w}} \log \tilde{p}(\mathbf{w}; \boldsymbol{\theta}),$$

which makes the Fisher divergence suitable for learning EBMs.

Score matching To get rid of the unknown $\nabla_{\mathbf{w}} \log p_D(\mathbf{w})$ in the Fisher divergence, Hyvärinen [24] proposes an equivalent form, named score matching (SM), as follows:

$$\mathcal{J}_{SM}(\boldsymbol{\theta}) \triangleq \mathbb{E}_{p_D(\mathbf{w})} \left[\frac{1}{2} ||\nabla_{\mathbf{w}} \log \tilde{p}(\mathbf{w}; \boldsymbol{\theta})||_2^2 + \text{tr}(\nabla_{\mathbf{w}}^2 \log \tilde{p}(\mathbf{w}; \boldsymbol{\theta})) \right] \equiv \mathcal{D}_F(p_D(\mathbf{w}) || p(\mathbf{w}; \boldsymbol{\theta})), \quad (2)$$

where $\nabla_{\mathbf{w}}^2 \log \tilde{p}(\mathbf{w}; \boldsymbol{\theta})$ is the Hessian matrix, $\text{tr}(\cdot)$ is the trace of a given matrix and \equiv means equivalence in parameter optimization. Though elegant, a straightforward application of SM is inefficient, as the computation of $\text{tr}(\nabla_{\mathbf{w}}^2 \log \tilde{p}(\mathbf{w}; \boldsymbol{\theta}))$ is time-consuming on high-dimensional data.

Sliced score matching To scale up SM, Song et al. [54] propose sliced score matching (SSM):

$$\mathcal{J}_{SSM}(\boldsymbol{\theta}) \triangleq \frac{1}{2} \mathbb{E}_{p_D(\mathbf{w})} [\|\nabla_{\mathbf{w}} \log \tilde{p}(\mathbf{w}; \boldsymbol{\theta})\|_2^2] + \mathbb{E}_{p_D(\mathbf{w})} \mathbb{E}_{p(\mathbf{u})} [\mathbf{u}^\top \nabla_{\mathbf{w}}^2 \log \tilde{p}(\mathbf{w}; \boldsymbol{\theta}) \mathbf{u}], \quad (3)$$

where \mathbf{u} is a random variable that is independent of \mathbf{w} and $p(\mathbf{u})$ satisfies certain mild conditions [54] to ensure that SSM is consistent with SM. Instead of calculating the trace of the Hessian matrix in SM, SSM computes the product of the Hessian matrix and a vector, which can be efficiently implemented by taking two normal back-propagation processes.

Denoising score matching Denoising score matching (DSM) [59] is another fast variant of SM:

$$\mathcal{J}_{DSM}(\boldsymbol{\theta}) \triangleq \mathbb{E}_{p_D(\mathbf{w}) p_\sigma(\tilde{\mathbf{w}}|\mathbf{w})} \|\nabla_{\tilde{\mathbf{w}}} \log \tilde{p}(\tilde{\mathbf{w}}; \boldsymbol{\theta}) - \nabla_{\tilde{\mathbf{w}}} \log p_\sigma(\tilde{\mathbf{w}}|\mathbf{w})\|_2^2 \equiv \mathcal{D}_F(p_\sigma(\tilde{\mathbf{w}}) \| p(\tilde{\mathbf{w}}; \boldsymbol{\theta})), \quad (4)$$

where $\tilde{\mathbf{w}}$ is the data perturbed by a noise distribution $p_\sigma(\tilde{\mathbf{w}}|\mathbf{w})$ with a hyperparameter σ and $p_\sigma(\tilde{\mathbf{w}}) = \int p_D(\mathbf{w}) p_\sigma(\tilde{\mathbf{w}}|\mathbf{w}) d\mathbf{w}$. A commonly chosen perturbation distribution is the Gaussian one that $p_\sigma(\tilde{\mathbf{w}}|\mathbf{w}) = \mathcal{N}(\tilde{\mathbf{w}}|\mathbf{w}, \sigma^2 I)$. DSM optimizes $\mathcal{D}_F(p_\sigma(\tilde{\mathbf{w}}) \| p(\tilde{\mathbf{w}}; \boldsymbol{\theta}))$ and is slightly inconsistent.

Multiscale denoising score matching Recently, Li et al. [38] propose multiscale denoising score matching (MDSM) to leverage different levels of noise to train EBMs on high-dimensional data as:

$$\mathcal{J}_{MDSM}(\boldsymbol{\theta}) \triangleq \mathbb{E}_{p_D(\mathbf{w}) p(\sigma) p_\sigma(\tilde{\mathbf{w}}|\mathbf{w})} \|\nabla_{\tilde{\mathbf{w}}} \log \tilde{p}(\tilde{\mathbf{w}}; \boldsymbol{\theta}) - \nabla_{\tilde{\mathbf{w}}} \log p_{\sigma_0}(\tilde{\mathbf{w}}|\mathbf{w})\|_2^2, \quad (5)$$

where $p(\sigma)$ is a prior distribution over the noise levels and σ_0 is a fixed noise level.

3 Method

In this paper, we aim to extend the above SM methods to learn general energy-based latent variable models (EBLVMs). In contrast to previous work [56, 58], our method only accesses the energy function without any structural assumption of the model. Formally, an EBLVM defines a probability distribution over a set of continuous visible variables \mathbf{v} and a set of latent variables \mathbf{h} as follows:

$$p(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta}) = \tilde{p}(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta}) / \mathcal{Z}(\boldsymbol{\theta}) = e^{-\mathcal{E}(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta})} / \mathcal{Z}(\boldsymbol{\theta}), \quad (6)$$

where $\mathcal{E}(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta})$ is the associated energy function with learnable parameters $\boldsymbol{\theta}$, $\tilde{p}(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta})$ is the unnormalized density, and $\mathcal{Z}(\boldsymbol{\theta}) = \int e^{-\mathcal{E}(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta})} d\mathbf{v} d\mathbf{h}$ is the partition function. In general, the marginal distribution $p(\mathbf{v}; \boldsymbol{\theta})$ and the posterior distribution $p(\mathbf{h}|\mathbf{v}; \boldsymbol{\theta})$ are intractable.

We would like to minimize $\mathcal{D}_F(q(\mathbf{v}) \| p(\mathbf{v}; \boldsymbol{\theta}))$, namely, the Fisher divergence between the marginal model distribution $p(\mathbf{v}; \boldsymbol{\theta})$ and $q(\mathbf{v})$, which can be the empirical data distribution $p_D(\mathbf{v})$ as in SM [24] or the perturbed one $p_\sigma(\tilde{\mathbf{v}}) = \int p_D(\mathbf{v}) p_\sigma(\tilde{\mathbf{v}}|\mathbf{v}) d\mathbf{v}$ in DSM [59]. Equivalently, we can optimize a certain SM objective in Eqn. (2-5), which is generally expressed in the following form:

$$\mathcal{J}(\boldsymbol{\theta}) = \mathbb{E}_{q(\mathbf{v}, \boldsymbol{\epsilon})} \mathcal{F}(\nabla_{\mathbf{v}} \log p(\mathbf{v}; \boldsymbol{\theta}), \boldsymbol{\epsilon}, \mathbf{v}), \quad (7)$$

where \mathcal{F} is a functional that depends on which SM objective we choose, $\boldsymbol{\epsilon}$ is introduced to represent additional random noise used in SSM [54] or DSM [59], and $q(\mathbf{v}, \boldsymbol{\epsilon})$ denotes the joint distribution of \mathbf{v} and $\boldsymbol{\epsilon}$. The same challenge for all SM objectives is that the marginal score function $\nabla_{\mathbf{v}} \log p(\mathbf{v}; \boldsymbol{\theta})$ is intractable and we propose bi-level score matching (BiSM) to solve the problem in this paper.

3.1 Bi-level Score Matching

First, we notice that the marginal score function can be rewritten as $\nabla_{\mathbf{v}} \log p(\mathbf{v}; \boldsymbol{\theta}) = \nabla_{\mathbf{v}} \log \frac{\tilde{p}(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta})}{p(\mathbf{h}|\mathbf{v}; \boldsymbol{\theta})} - \nabla_{\mathbf{v}} \log \mathcal{Z}(\boldsymbol{\theta}) = \nabla_{\mathbf{v}} \log \frac{\tilde{p}(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta})}{p(\mathbf{h}|\mathbf{v}; \boldsymbol{\theta})}$. We introduce a variational posterior distribution $q(\mathbf{h}|\mathbf{v}; \boldsymbol{\phi})$ to approximate the true posterior $p(\mathbf{h}|\mathbf{v}; \boldsymbol{\theta})$, and obtain an approximation of the marginal score function using $\nabla_{\mathbf{v}} \log \frac{\tilde{p}(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta})}{q(\mathbf{h}|\mathbf{v}; \boldsymbol{\phi})}$. We reformulate the general SM objective in Eqn. (7) as the following bi-level optimization problem:

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta} \in \Theta} \mathcal{J}_{Bi}(\boldsymbol{\theta}, \boldsymbol{\phi}^*(\boldsymbol{\theta})), \quad \mathcal{J}_{Bi}(\boldsymbol{\theta}, \boldsymbol{\phi}) = \mathbb{E}_{q(\mathbf{v}, \boldsymbol{\epsilon})} \mathbb{E}_{q(\mathbf{h}|\mathbf{v}; \boldsymbol{\phi})} \mathcal{F} \left(\nabla_{\mathbf{v}} \log \frac{\tilde{p}(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta})}{q(\mathbf{h}|\mathbf{v}; \boldsymbol{\phi})}, \boldsymbol{\epsilon}, \mathbf{v} \right), \quad (8)$$

²See Sec. 4 for possible extensions on discrete \mathbf{v} .

where Θ is the hypothesis space of the model and $\phi^*(\theta)$ is defined as follows:

$$\phi^*(\theta) = \arg \min_{\phi \in \Phi} \mathcal{G}(\theta, \phi), \text{ with } \mathcal{G}(\theta, \phi) = \mathbb{E}_{q(\mathbf{v}, \epsilon)} \mathcal{D}(q(\mathbf{h}|\mathbf{v}; \phi) || p(\mathbf{h}|\mathbf{v}; \theta)), \quad (9)$$

where Φ is the hypothesis space of the variational posterior and \mathcal{D} is a certain divergence to be specified later. We denote ϕ^* as a function of θ to explicitly present the dependency. We emphasize that the bi-level formulation is necessary because if we treat ϕ^* as a constant with respect to θ , then $\nabla_{\theta} \mathcal{J}(\theta) \neq \nabla_{\theta} \mathcal{J}_{Bi}(\theta, \phi^*)$ in general. In contrast, the equivalence of the bi-level formulation in Eqn. (8-9) and SM under the nonparametric assumption [15] is characterized in Theorem 1.

Theorem 1. (Equivalence of BiSM, proof in Appendix A.1) Assuming that $\forall \theta \in \Theta, \exists \phi \in \Phi$ such that $\mathcal{D}(q(\mathbf{h}|\mathbf{v}; \phi) || p(\mathbf{h}|\mathbf{v}; \theta)) = 0, \forall \mathbf{v} \in \text{supp}(q)$, we have $\nabla_{\theta} \mathcal{J}(\theta) = \nabla_{\theta} \mathcal{J}_{Bi}(\theta, \phi^*(\theta))$.

We notice that such assumptions have been made in the typical variational inference [29] to obtain a tight estimate of the log-likelihood and the generative adversarial networks [15] to guarantee validity. Further, in a practical case with a less powerful $q(\mathbf{h}|\mathbf{v}; \phi)$, we bound the bias of BiSM by the approximation error of $q(\mathbf{h}|\mathbf{v}; \phi)$ in Appendix A.1 to provide a complementary analysis.

To learn general EBLVMs with intractable posteriors, the lower level optimization problem in Eqn. (9) can only access the unnormalized model distribution $\tilde{p}(\mathbf{v}, \mathbf{h}; \theta)$ and the variational posterior $q(\mathbf{h}|\mathbf{v}; \phi)$ in calculation, as in the high-level problem in Eqn. (8). We first consider the widely adopted KL divergence for variational inference and obtain an equivalent form regarding to optimizing ϕ :

$$\mathcal{D}_{KL}(q(\mathbf{h}|\mathbf{v}; \phi) || p(\mathbf{h}|\mathbf{v}; \theta)) \equiv \mathbb{E}_{q(\mathbf{h}|\mathbf{v}; \phi)} \log \frac{q(\mathbf{h}|\mathbf{v}; \phi)}{\tilde{p}(\mathbf{v}, \mathbf{h}; \theta)}, \quad (10)$$

from which an unknown constant is subtracted. Therefore, Eqn. (10) is sufficient for training ϕ but not suitable for evaluating the inference accuracy. In contrast, the Fisher divergence, which is an alternative approach for variational inference [62], can be directly calculated by:

$$\mathcal{D}_F(q(\mathbf{h}|\mathbf{v}; \phi) || p(\mathbf{h}|\mathbf{v}; \theta)) = \frac{1}{2} \mathbb{E}_{q(\mathbf{h}|\mathbf{v}; \phi)} [\|\nabla_{\mathbf{h}} \log q(\mathbf{h}|\mathbf{v}; \phi) - \nabla_{\mathbf{h}} \log \tilde{p}(\mathbf{v}, \mathbf{h}; \theta)\|_2^2]. \quad (11)$$

For the detailed derivation, please see Appendix A.2. Compared with the KL divergence in Eqn. (10), the Fisher divergence in Eqn. (11) can be used for both training and evaluation but cannot deal with discrete \mathbf{h} in which case $\nabla_{\mathbf{h}}$ is not well defined. In our experiments, we apply them according to the specific scenario. In principle, any other divergence that does not necessarily access $p(\mathbf{v}; \theta)$ or $p(\mathbf{h}|\mathbf{v}; \theta)$ can be used here and we leave a systematical study of the divergence for the future work.

3.2 Stochastic Optimization for BiSM

Our goal is to learn general EBLVMs whose energy can be parameterized by highly nonlinear and nonstructural functions, e.g. deep neural networks (DNNs). Such models are rarely studied before due to the challenges in training and inference. In Sec. 5.2, we propose an instance motivated by the fully visible EBMs [9, 38] to fit natural images and validate the effectiveness of BiSM. To fit the intractable posterior, we also employ a variational posterior parameterized by DNNs. In this context, it is impractical to exactly solve the BiSM problem in Eqn. (8-9) on full batch. Therefore, we develop a practical algorithm for BiSM by updating ϕ and θ alternatively using stochastic gradient descent. Formally, assuming a minibatch of data is given, let $\hat{\mathcal{J}}(\theta)$, $\hat{\mathcal{J}}_{Bi}(\theta, \phi)$ and $\hat{\mathcal{G}}(\theta, \phi)$ denote the corresponding functions evaluated on the minibatch. Motivated by our analysis in Theorem 2 presented later, we first update ϕ for K times on the same minibatch of data by:

$$\phi \leftarrow \phi - \alpha \frac{\partial \hat{\mathcal{G}}(\theta, \phi)}{\partial \phi}, \quad (12)$$

where α is a prefixed learning rate scheme. We denote the resulting parameters of the variational posterior as ϕ^0 , which approximates $\hat{\phi}^*(\theta) = \arg \min_{\phi \in \Phi} \hat{\mathcal{G}}(\theta, \phi)$. To update θ , the central challenge is to approximate the stochastic gradient $\frac{\partial \hat{\mathcal{J}}_{Bi}(\theta, \hat{\phi}^*(\theta))}{\partial \theta}$, which is addressed by the gradient unrolling technique [41]. Specifically, we start from ϕ^0 and calculate $\hat{\phi}^N(\theta)$ recursively by:

$$\hat{\phi}^1(\theta) = \phi^0 - \alpha \frac{\partial \hat{\mathcal{G}}(\theta, \phi)}{\partial \phi} \Big|_{\phi=\phi^0}, \text{ and } \hat{\phi}^n(\theta) = \hat{\phi}^{n-1}(\theta) - \alpha \frac{\partial \hat{\mathcal{G}}(\theta, \phi)}{\partial \phi} \Big|_{\phi=\hat{\phi}^{n-1}(\theta)}, \quad (13)$$

for $n = 2, \dots, N$, where we treat ϕ^0 as a constant with respect to θ , and make the dependence of $\hat{\phi}^n(\theta)$ on ϕ^0 implicit for simplicity. Note that Eqn. (13) is not used to update the variational parameters but to approximate $\frac{\partial \hat{\mathcal{J}}_{Bi}(\theta, \hat{\phi}^*(\theta))}{\partial \theta}$ by the gradient of a surrogate loss $\hat{\mathcal{J}}_{Bi}(\theta, \hat{\phi}^N(\theta))$:

$$\frac{\partial \hat{\mathcal{J}}_{Bi}(\theta, \hat{\phi}^N(\theta))}{\partial \theta} = \frac{\partial \hat{\mathcal{J}}_{Bi}(\theta, \phi)}{\partial \theta} \Big|_{\phi=\hat{\phi}^N(\theta)} + \frac{\partial \hat{\mathcal{J}}_{Bi}(\theta, \phi)}{\partial \phi} \Big|_{\phi=\hat{\phi}^N(\theta)} \frac{\partial \hat{\phi}^N(\theta)}{\partial \theta}. \quad (14)$$

Given the above stochastic gradient, we update the parameters in the model distribution by:

$$\theta \leftarrow \theta - \beta \frac{\partial \hat{\mathcal{J}}_{Bi}(\theta, \hat{\phi}^N(\theta))}{\partial \theta}, \quad (15)$$

where β is a prefixed learning rate scheme. The whole training procedure is summarized in Algorithm 1. The gradient unrolling technique has been proposed to learn implicit directed generative models [41, 9]. In comparison, we learn EBMs and construct a different bi-level optimization problem based on the Fisher divergence. Further, below we formally analyze the approximation error of the stochastic gradient, which has not been explored in the related work [41, 9].

Theorem 2. (Approximate error of the stochastic gradient, proof in Appendix A.4) *Supposing that:*

1. Both Θ and Φ are compact and convex,
2. $\hat{\mathcal{J}}_{Bi}(\theta, \phi) \in C^2(\Omega)$, $\hat{\mathcal{G}}(\theta, \phi) \in C^3(\Omega)$, where Ω is an open set including $\Theta \times \Phi$ (i.e. $\hat{\mathcal{J}}_{Bi}$ and $\hat{\mathcal{G}}$ are second and third order continuously differentiable on Ω respectively),
3. $\hat{\mathcal{G}}(\theta, \phi)$ is strongly convex on Φ for all $\theta \in \Theta$,
4. $\forall N \geq 0, \forall \theta \in \Theta, \forall \phi^0 \in \Phi, \hat{\phi}^N(\theta, \phi^0) \in \Phi$ and $\hat{\phi}^*(\theta) \in \Phi$,

then when α is small enough, there exists $A, B, C > 0$ and $\kappa \in (0, 1)$ independent of θ and ϕ^0 , s.t.,

$$\left\| \frac{\partial \hat{\mathcal{J}}_{Bi}(\theta, \hat{\phi}^N(\theta, \phi^0))}{\partial \theta} - \frac{\partial \hat{\mathcal{J}}_{Bi}(\theta, \hat{\phi}^*(\theta))}{\partial \theta} \right\| \leq (A + BN)\kappa^N \|\phi^0 - \hat{\phi}^*(\theta)\| + C\kappa^N,$$

for all $\theta \in \Theta, \phi^0 \in \Phi$ and $N \geq 0$.

Theorem 2 implies that the approximation error converges to zero in a linear rate in terms of N when $\hat{\mathcal{G}}$ is strongly convex, which is a commonly used assumption to obtain such results [3]. Although the assumption does not always hold in our experiments, Theorem 2 still provides insights into our implementation. As mentioned before, we update ϕ for K times on the same minibatch of data to reduce $\|\phi^0 - \hat{\phi}^*\|$. In Fig. 1, we numerically validate Theorem 2. The gradient bias decays (approximately) exponentially w.r.t. N , which is consistent with Theorem 2. As for $\|\phi^0 - \hat{\phi}^*\|$, we find it decreases from 1.38 to 0.87 as K increases from 0 to 20. It leads to smaller bias, which agrees with Fig. 1 and Theorem 2. Besides, we notice that the unrolling technique is not exclusive for advances in non-convex optimization [3] and we leave the general analysis for the future work. Due to the controllable approximation error, the convergence of BiSM can be formally characterized as follows.

Corollary 3. (BiSM finds δ -stationary points, proof in Appendix A.5) *For any accuracy level $\delta > 0$, assuming Theorem 2 holds, using a sufficiently large N , i.e. asymptotically $\mathcal{O}(\log \frac{1}{\delta})$, and a proper learning rate scheme β [3], Algorithm 1 converges to a δ -stationary point of BiSM in Eqn. (8-9), and further a δ -stationary point of SM in Eqn. (7) if Theorem 1 also holds.*

4 Related Work

Apart from the MLE and SM-based methods mentioned before, we present other related work on learning and inference in energy-based latent variable models (EBLVMs) and bi-level optimization.




Figure 1: The gradient bias (the left hand side of Theorem 2) w.r.t. N and K in GRBM on Frey face.

Previous state-of-the-art deep EBLVMs include deep belief networks (DBNs) [23], convolutional deep belief networks (CDBNs) [36] and deep Boltzmann machines (DBMs) [48]. In comparison, first, their unsupervised pretraining algorithms explicitly leverage the layer-wise structures in their models. In contrast, BiSM is an end-to-end approach which does not require any specific structural assumption. A direct comparison of BiSM and the layer-wise algorithms in DBNs and DBMs is nontrivial because the model likelihoods are discrete and then the Fisher divergence is not well-defined. We mention that extensions of SM [25, 39, 52] can deal with discrete data and BiSM can be applied to such methods as well but we leave it for the future work. Second, these methods either focus on supervised learning [36] or model relatively simple data [23, 48, 49]. Our experiments show a deep EBLVM trained by BiSM can successfully generate natural images in a purely unsupervised learning setting, which has not been explored before. Besides, it is possible to generalize BiSM to (semi-)supervised learning like the extensions [33, 50, 6] of directed deep generative models [31, 15].

Noise contrastive estimation (NCE) [17] is another criterion to learn unnormalized models. It discriminates samples from the model distribution and a prefixed noise distribution. The main bottleneck of NCE on high-dimensional data is how to define a proper noise distribution manually. Variational noise-contrastive estimation [47] applies variational inference to the NCE objective and therefore can be used for posterior inference of latent variables. However, the difficulty of choosing the noise distribution remains and only small experiments are reported in VNCE [47].

The bi-level optimization problem arises in different tasks, such as learning implicit directed generative models [15, 41, 9], meta learning [12, 46, 13], and many others [40, 8, 1, 2]. Franceschi et al. [13] provide some other properties of gradient unrolling [41], from which a result similar to Corollary 3 is derived. As an alternative approach to gradient unrolling [41], implicit gradient formulation is adopted in [8, 1, 46, 2]. In comparison, implicit gradient can be accurate in a non-asymptotic setting but needs to solve a quadratic problem involving the inverse of the Hessian matrix w.r.t. the parameters in the lower level problem. Our preliminary experiments suggest that it is impractical to use the implicit gradient in BiSM if the variational posterior is a deep convolutional neural network.

5 Experiment

We evaluate BiSM on two EBLVMs in our experiments. The first one is the well-known Gaussian restricted Boltzmann machine (GRBM) [60, 22], which is a good benchmark to compare with existing methods [22, 24]. The second one is a deep EBLVM introduced in this paper to model natural image data. We would like to validate two arguments: (1) BiSM is comparable to strong baselines, including the contrastive divergence (CD)-based methods [21, 57] and SM-based methods [56, 54], when they are applicable; and (2) BiSM can learn general EBLVMs with neural network energy functions and generate natural images like CIFAR10, which have not been explored in previous work to our knowledge. We explicitly denote our method as *BiDSM*, *BiSSM* and *BiMDSM*, respectively, based on which SM objective [59, 54, 38] is used. All of them are collectively called *BiSM* for simplicity.

5.1 GRBM

The energy function of a GRBM is $\mathcal{E}(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta}) = \frac{1}{2\sigma^2} \|\mathbf{v} - b\|^2 - c^\top \mathbf{h} - \frac{1}{\sigma} \mathbf{v}^\top W \mathbf{h}$, where the learnable parameters are $\boldsymbol{\theta} = (\sigma, W, b, c)$. To deal with discrete \mathbf{h} , we use the KL divergence in Eqn. (10) and the Gumbel-Softmax trick [27] in the lower-level problem. Because $p(\mathbf{h}|\mathbf{v}; \boldsymbol{\theta})$ is tractable, DSM [11] and SSM [54] can learn a GRBM as the way described in [56], which serve as gold standard baselines of our method. Besides, we also compare BiSM against the popular CD-based methods [21, 57] and the NCE-based methods [17, 47].

Algorithm 1 Bi-level score matching by alternative stochastic gradient descent

- 1: **Input:** Constants K and N , learning rate schemes α and β , randomly initialized $\boldsymbol{\theta}$ and ϕ
 - 2: **repeat**
 - 3: Sample a minibatch of data
 - 4: **for** $i = 1, \dots, K$ **do**
 - 5: Update ϕ according to Eqn. (12)
 - 6: **end for**
 - 7: $\phi^0 \leftarrow \phi$
 - 8: **for** $n = 1, \dots, N$ **do**
 - 9: Compute $\hat{\phi}^N(\boldsymbol{\theta}, \phi^0)$ according to Eqn. (13)
 - 10: **end for**
 - 11: Approximate the stochastic gradient according to Eqn. (14) and update $\boldsymbol{\theta}$ according to Eqn. (15)
 - 12: **until** Convergence or reaching certain threshold
-




Figure 2: A small GRBM trained by different methods to fit the checkerboard dataset. Our BiSSM-5 and BiDSM-5 achieve comparable performance to the SM, CD and NCE baselines.




Figure 3: Test log-likelihood on the checkerboard dataset (averaged over 10 runs).

Table 1: Test log-likelihood (LL) and test Fisher divergence (Fisher) (subtracted by the same unknown constant only relevant to the data) according to the best validation performance on the Frey face dataset.

Method	LL \uparrow	Fisher \downarrow
DSM	129.23	-5885.09
BiDSM ($N=0$)	107.59	-5474.52
BiDSM ($N=1$)	110.65	-5516.07
BiDSM ($N=5$)	124.00	-5780.18
BiDSM ($N=10$)	125.72	-5800.17

We briefly summarize the default experimental settings here, see Appendix B.1 and the source code³ for more details. We evaluate BiSM on the checkerboard dataset and the Frey face dataset. The checkerboard dataset consists of 2-D points and the density is shown in Fig. 2 (a). We generate 60,000 points for training and 10,000 points for testing. The dimension of \mathbf{h} is 4. The Frey face dataset consists of gray-scaled face images of size 20×28 . We split 1,400 images for training, 300 images for validation and 265 images for testing. The dimension of \mathbf{h} is 400. On both datasets, $q(\mathbf{h}|v; \phi)$ is a Bernoulli distribution parameterized by a fully connected layer with the sigmoid activation and we use $K = 5$ and $N = 5$ in BiSM. We use the Adam [30] optimizer for all methods. The learning rate is 10^{-3} on the checkerboard dataset and 2×10^{-4} on the Frey face dataset.

In Fig. 2, we plot the density of the same GRBM trained by different methods on the checkerboard dataset. Our BiSSM and BiDSM and the corresponding SM baselines [54, 59] are comparable after convergence as shown in Fig. 2 (e-h), demonstrating that BiSM can extend SM to deal with EBLVMs in a black-box manner without hurting performance. The results of CD-based methods in Fig. 2 (b-d) and NCE-based methods in Fig. 2 (i-j) are similar to BiSM. Besides, we calculate the test log-likelihood by brute force in Fig. 3 for a quantitative comparison. The results after convergence agree with the results in Fig. 2. The convergence speed of BiSSM is as fast as SSM while that of BiDSM is slightly slower than DSM.

We provide the results on the Frey face dataset in Tab. 1. We consider the log-likelihood as well as the score matching loss, which is the Fisher divergence subtracted by an unknown constant that is only relevant to the data distribution. In this case, SSM [54] needs additive noise on the data [53] and is equivalent to DSM. CD-based methods require the model and data distribution to be properly modified [5, 4, 37], resulting in incomparable (unnormalized) Fisher divergence. Therefore, we focus

³<https://github.com/baofff/BiSM>

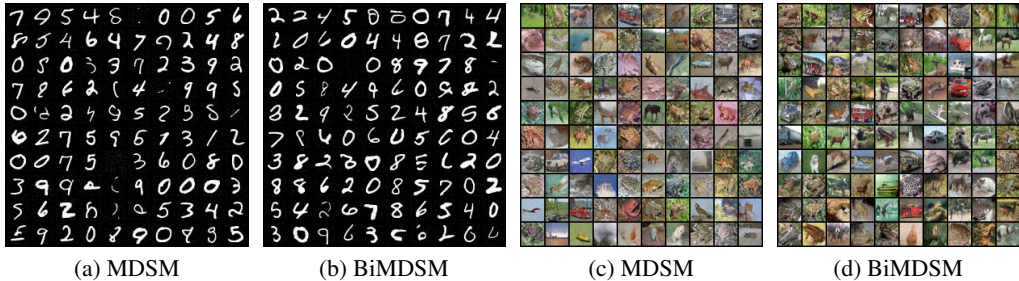


Figure 4: Samples from EBMs trained by MDSM and comparable EBLVMs trained by BiMDSM on the MNIST and CIFAR10 datasets. The samples from both models are of similar visual quality.




Figure 5: Learning curves of BiSM on MNIST and CIFAR10. Fig. (a-b) show the learning curve of $\hat{\mathcal{J}}_{Bi}(\theta, \hat{\phi}^N(\theta))$. Fig. (c-d) show the learning curve of $\mathcal{D}_F(q(\mathbf{h}|\mathbf{v})||p(\mathbf{h}|\mathbf{v}))$.

on the comparison with DSM here. According to Tab. 1, as N increases, BiDSM gets better due to the smaller approximation error of the stochastic gradient and BiDSM is comparable to DSM with $N \geq 5$. See Appendix C.1 for the samples, additional sensitivity analysis and time complexity comparison.

5.2 Deep EBLVM

We introduce a deep EBLVM with energy function $\mathcal{E}(\mathbf{v}, \mathbf{h}; \theta) = g_3(g_2(g_1(\mathbf{v}; \theta_1), \mathbf{h}); \theta_2)$, where $\theta = (\theta_1, \theta_2)$. $g_1(\cdot)$ is a neural network that outputs a feature sharing the same dimension with \mathbf{h} and the architecture is adopted from the deep EBM proposed in MDSM [38]. $g_2(\cdot, \cdot)$ is an additive coupling layer [7] to make the features and the latent variables strongly coupled. $g_3(\cdot)$ is a small neural network that outputs a scalar. Because the posterior is intractable in the deep EBLVM, the baselines used in Sec.5.1 are not applicable. Further, to our knowledge, such models have not been explored in related work on learning deep EBLVMs [49, 34, 37] due to their limitations discussed before. Therefore, we use a fully visible EBM trained by MDSM [38] as our baseline. We emphasize that our goal is not to achieve state-of-the-art results on the related tasks but to validate that BiSM can learn complex EBLVMs to generate natural images and perform posterior inference accurately.

We evaluate both methods on the MNIST and CIFAR10 datasets. The MNIST dataset consists of gray-scaled hand-written digits of size 28×28 and the CIFAR10 dataset consists of color natural images of size 32×32 . Following MDSM [38], we split 60,000 samples for training and 10,000 samples for testing on MNIST, and split 50,000 samples for training and 10,000 samples for testing on CIFAR10; we use a 12-layer ResNet [19] and a 18-layer ResNet in $g_1(\cdot)$ on the MNIST and CIFAR10 datasets respectively. On both datasets, we use a fully connected layer in $g_3(\cdot)$. Overall, our EBLVMs have comparable parameters to the models of the baseline [38]. On both datasets, the dimension of \mathbf{h} is 50 by default, and $q(\mathbf{h}|\mathbf{v}; \phi)$ is a Gaussian distribution parameterized by a 3-layer convolutional neural network (CNN). We use $K = 5$ and $N = 0$ for time and memory efficiency and the Adam [30] optimizer with learning rates 10^{-4} and 5×10^{-5} for training on MNIST and CIFAR10 respectively following [38]. See further details (e.g., the sampling procedure) in Appendix B.2.

In Fig.4, we present the samples on the MNIST and CIFAR10 datasets. It can be seen that our method produces samples of similar visual quality to the baseline’s, indicating that the deep EBLVM learned by BiMDSM captures the marginal data distribution. Tab. 2 shows the FID [20] scores of existing EBMs. Following a similar protocol adopted in MDSM [38], the third column shows the FID with early stopping (ES) according to the results on 1,000 samples. We also implement MDSM in our

code for a fair comparison. The reproduced MDSM is slightly better than the original paper [38] and serves as a stronger baseline. Our result outperforms MDSM in a fair comparison, indicating that introducing latent variables can improve the sample quality. We mention that Flow-CE [14] and VAE-EBLVM [18] generate samples from a flow and a VAE respectively instead of EBMs are less comparable. We include them because such results have not been reported in previous state-of-the-art EBLVMs [23, 49, 36]. We hope our results can serve as a benchmark for the future work on EBLVMs.

As for posterior inference (the lower level problem), Fig. 5 (c-d) shows the learning curves of the Fisher divergence (see Eqn. (11)) between the variational posterior $q(\mathbf{h}|\mathbf{v})$ and the model posterior $p(\mathbf{h}|\mathbf{v})$ with different dimensions of \mathbf{h} on the MNIST and CIFAR10 datasets. In all settings, the divergence decreases from a magnitude of 10^5 or 10^4 to a magnitude of 10^2 or 10^1 , suggesting that $q(\mathbf{h}|\mathbf{v})$ learned by BiSM is an accurate approximation of $p(\mathbf{h}|\mathbf{v})$. On CIFAR10, the divergence sometimes increases perhaps because $p(\mathbf{h}|\mathbf{v})$ gets complex during training. Nevertheless, the divergence after convergence is still relatively small. We also show the learning curves of the higher level problem in Fig. 5 (a-b), which decreases stably. We provide results on conditionally sampling, feature embedding and classification in Appendix C.2.

Table 2: FID on CIFAR10. \dagger means averaged by 5 runs. Methods with \ddagger use comparable networks.

Method	FID \downarrow	FID-ES \downarrow
Flow-CE [14]	37.30	-
VAE-EBLVM [18]	30.1	-
CoopNets [61]	33.61	-
EBM (ensemble) [10]	38.2	-
MDSM \ddagger [38]	-	31.7
MDSM \ddagger (our code)	39.12	$30.19 \pm 2.60^\dagger$
BiMDSM \ddagger (20)	34.55	$26.62 \pm 1.52^\dagger$
BiMDSM \ddagger (50)	38.82	$29.43 \pm 2.76^\dagger$
BiMDSM \ddagger (100)	36.13	$26.90 \pm 2.14^\dagger$

6 Conclusion and Discussion

We consider to extend score matching (SM) to learn energy-based latent variable models with a minimal model assumption. We reformulate a SM objective as a bi-level optimization problem, named bi-level score matching (BiSM), which introduces a variational distribution for posterior inference. We prove the equivalence of BiSM to SM under the nonparametric assumption. We develop an efficient stochastic optimization algorithm with gradient unrolling for BiSM and provide a formal convergence analysis. We show the promise of BiSM in Gaussian restricted Boltzmann machines and a highly nonstructural EBLVM parameterized by deep neural networks. BiSM is comparable to the widely adopted contrastive divergence and SM methods when they are applicable; and can learn complex EBLVMs with intractable posteriors to generate natural images.

Though promising, the bilevel optimization involved in BiSM is complicated and time-consuming. A potential future work is to simplify BiSM by directly approximating the gradient of the score function with respect to the model parameters in order to avoid such bilevel formulation.

Broader Impact

In many real world applications, such as bioinformatics, social network analysis and so on, sometimes energy-based models are preferable than directed models. The ability of the proposed BiSM to learn general energy-based latent variable models can potentially benefit such applications and therefore benefit the society.

However, as a way to train deep generative models, this work can be abused to produce fake images, videos and news, similarly to the generative adversarial nets.

Acknowledgments and Disclosure of Funding

We thank Ziyu Wang, Yucen Luo, Tianyu Pang and Yang Song for feedback on our work, and we thank Cheng Lu, Yuhao Zhou and Shihong Song for proofreading. This work was supported by the National Key Research and Development Program of China (Nos. 2017YFA0700904, 2020AAA0104304), NSFC Projects (Nos. 61620106010, 62076145, 62076147, U19B2034, U1811461, U19A2081), Beijing NSF Project (No. L172037), Beijing Academy of Artificial Intelligence (BAAI), THU-Bosch JCMML center, Tsinghua-Huawei Joint Research Program, a grant from Tsinghua Institute for Guo Qiang, Tiangong Institute for Intelligent Computing, the JP Morgan Faculty Research Program and the NVIDIA NVAIL Program with GPU/DGX Acceleration. C. Li was supported by the Chinese postdoctoral innovative talent support program and Shuimu Tsinghua Scholar.

References

- [1] Brandon Amos and J Zico Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 136–145. JMLR. org, 2017.
- [2] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems*, pages 688–699, 2019.
- [3] Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- [4] Kyung Hyun Cho, Tapani Raiko, and Alexander Ilin. Gaussian-bernoulli deep boltzmann machine. In *Neural Networks (IJCNN), The 2013 International Joint Conference on*, pages 1–7. IEEE, 2013.
- [5] KyungHyun Cho, Alexander Ilin, and Tapani Raiko. Improved learning of gaussian-bernoulli restricted boltzmann machines. In *International conference on artificial neural networks*, pages 10–17. Springer, 2011.
- [6] LI Chongxuan, Taufik Xu, Jun Zhu, and Bo Zhang. Triple generative adversarial nets. In *Advances in neural information processing systems*, pages 4088–4098, 2017.
- [7] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- [8] Justin Domke. Generic methods for optimization-based modeling. In *Artificial Intelligence and Statistics*, pages 318–326, 2012.
- [9] Chao Du, Kun Xu, Chongxuan Li, Jun Zhu, and Bo Zhang. Learning implicit generative models by teaching explicit ones. *arXiv preprint arXiv:1807.03870*, 2018.
- [10] Yilun Du and Igor Mordatch. Implicit generation and generalization in energy-based models. *arXiv preprint arXiv:1903.08689*, 2019.
- [11] Chelsea Finn, Paul Christiano, Pieter Abbeel, and Sergey Levine. A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *arXiv preprint arXiv:1611.03852*, 2016.
- [12] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- [13] Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil. Bilevel programming for hyperparameter optimization and meta-learning. *arXiv preprint arXiv:1806.04910*, 2018.
- [14] Ruiqi Gao, Erik Nijkamp, Diederik P Kingma, Zhen Xu, Andrew M Dai, and Ying Nian Wu. Flow contrastive estimation of energy-based models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7518–7528, 2020.
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [16] Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. Your classifier is secretly an energy based model and you should treat it like one. *arXiv preprint arXiv:1912.03263*, 2019.
- [17] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304, 2010.

- [18] Tian Han, Erik Nijkamp, Linqi Zhou, Bo Pang, Song-Chun Zhu, and Ying Nian Wu. Joint training of variational auto-encoder and latent energy-based model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7978–7987, 2020.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
- [20] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017.
- [21] Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002.
- [22] Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *science*, 313(5786):504–507, 2006.
- [23] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006.
- [24] Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(Apr):695–709, 2005.
- [25] Aapo Hyvärinen. Some extensions of score matching. *Computational Statistics & Data Analysis*, 51(5):2499–2512, February 2007. URL <https://ideas.repec.org/a/eee/csdana/v51y2007i5p2499-2512.html>.
- [26] John Ingraham and Debora Marks. Variational inference for sparse and undirected models. In *International Conference on Machine Learning*, pages 1607–1616. PMLR, 2017.
- [27] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [28] Oliver Thomas Johnson. *Information theory and the central limit theorem*. World Scientific, 2004.
- [29] Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- [30] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [31] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [32] Durk P Kingma and Yann L Cun. Regularized estimation of image statistics by score matching. In *Advances in neural information processing systems*, pages 1126–1134, 2010.
- [33] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, pages 3581–3589, 2014.
- [34] Volodymyr Kuleshov and Stefano Ermon. Neural variational inference and learning in undirected graphical models. In *Advances in Neural Information Processing Systems*, pages 6734–6743, 2017.
- [35] Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and F Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
- [36] Honglak Lee, Roger Grosse, Rajesh Ranganath, and Andrew Y Ng. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *Proceedings of the 26th annual international conference on machine learning*, pages 609–616, 2009.

- [37] Chongxuan Li, Chao Du, Kun Xu, Max Welling, Jun Zhu, and Bo Zhang. To relieve your headache of training an mrf, take advil. *arXiv preprint arXiv:1901.08400*, 2019.
- [38] Zengyi Li, Yubei Chen, and Friedrich T Sommer. Annealed denoising score matching: Learning energy-based models in high-dimensional spaces. *arXiv preprint arXiv:1910.07762*, 2019.
- [39] Siwei Lyu. Interpretation and generalization of score matching. In *In Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, 2009.
- [40] Dougal Maclaurin, David Duvenaud, and Ryan P. Adams. Gradient-based hyperparameter optimization through reversible learning. In *Proceedings of the 32nd International Conference on Machine Learning - Volume 37, ICML'15*, page 2113–2122. JMLR.org, 2015.
- [41] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*, 2016.
- [42] Radford M. Neal. Probabilistic inference using markov chain monte carlo methods, 1993.
- [43] Erik Nijkamp, Mitch Hill, Tian Han, Song-Chun Zhu, and Ying Nian Wu. On the anatomy of mcmc-based maximum likelihood learning of energy-based models. *arXiv preprint arXiv:1903.12370*, 2019.
- [44] Erik Nijkamp, Song-Chun Zhu, and Ying Nian Wu. On learning non-convergent short-run mcmc toward energy-based model. *arXiv preprint arXiv:1904.09770*, 2019.
- [45] Tianyu Pang, Kun Xu, Chongxuan Li, Yang Song, Stefano Ermon, and Jun Zhu. Efficient learning of generative models via finite-difference score matching. *arXiv preprint arXiv:2007.03317*, 2020.
- [46] Aravind Rajeswaran, Chelsea Finn, Sham M. Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *NeurIPS*, 2019.
- [47] Benjamin Rhodes and Michael U. Gutmann. Variational noise-contrastive estimation. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 2741–2750. PMLR, 16–18 Apr 2019. URL <http://proceedings.mlr.press/v89/rhodes19a.html>.
- [48] Ruslan Salakhutdinov and Geoffrey Hinton. Deep Boltzmann machines. In *Proceedings of the twelfth international conference on artificial intelligence and statistics*, 2009.
- [49] Ruslan Salakhutdinov and Hugo Larochelle. Efficient learning of deep Boltzmann machines. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 693–700, 2010.
- [50] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
- [51] Saeed Saremi, Arash Mehrjou, Bernhard Schölkopf, and Aapo Hyvärinen. Deep energy estimator networks. *arXiv preprint arXiv:1805.08306*, 2018.
- [52] Jascha Sohl-Dickstein, Peter Battaglino, and Michael R DeWeese. Minimum probability flow learning. *arXiv preprint arXiv:0906.4779*, 2009.
- [53] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 11895–11907, 2019.
- [54] Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation. *arXiv preprint arXiv:1905.07088*, 2019.
- [55] Nitish Srivastava and Russ R Salakhutdinov. Multimodal learning with deep boltzmann machines. In *Advances in neural information processing systems*, pages 2222–2230, 2012.

- [56] Kevin Swersky, Marc’ Aurelio Ranzato, David Buchman, Nando D Freitas, and Benjamin M Marlin. On autoencoders and score matching for energy based models. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 1201–1208, 2011.
- [57] Tijmen Tieleman. Training restricted boltzmann machines using approximations to the likelihood gradient. In *Proceedings of the 25th international conference on Machine learning*, pages 1064–1071. ACM, 2008.
- [58] Eszter Vértes, UCL Gatsby Unit, and Maneesh Sahani. Learning doubly intractable latent variable models via score matching. <http://approximateinference.org/2016/accepted/VertesSahani2016.pdf>, 2016.
- [59] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [60] Max Welling, Michal Rosen-Zvi, and Geoffrey E Hinton. Exponential family harmoniums with an application to information retrieval. In *Advances in neural information processing systems*, pages 1481–1488, 2005.
- [61] Jianwen Xie, Yang Lu, Ruiqi Gao, Song-Chun Zhu, and Ying Nian Wu. Cooperative training of descriptor and generator networks. *IEEE transactions on pattern analysis and machine intelligence*, 42(1):27–45, 2018.
- [62] Yue Yang, Ryan Martin, and Howard Bondell. Variational approximations using fisher divergence. *arXiv preprint arXiv:1905.05284*, 2019.
- [63] Qiu Yixuan, Zhang Lingsong, and Wang Xiao. Unbiased contrastive divergence algorithm for training energy-based latent variable models. In *International Conference on Learning Representations*, 2020.